

10/03/00

10-04-00

A

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
REQUEST FOR FILING APPLICATION Under Rule 53(a),(b)&(f)
(No Filing Fee or Declaration); RULE 53(f) NO DECLARATION

PATENT
APPLICATION

Asst. Commissioner for Patents
Washington, DC 20231

Atty. Dkt.

83818/0269851

TEN-007

C/M#

Client Ref

Date: October 3, 2000

Sir:

1. This is a Request for filing a new utility PATENT APPLICATION entitled: **Debugging Apparatus and Method For Systems of Configurable Processors** without a filing fee or Oath/Declaration but for which is enclosed the following:

2. 35 pages of spec, claims and abstract;
3. 37 pages of Appendix A (page nos. 36-72)
4. 20 claims;
5. ☒ Drawings: 2 sheets of formal; 8½x11 sized paper (Figs. 1-3);
6. ☒ USPTO Form 1449 citing 6 references; and
7. ☒ Copies of the cited References

6. This application is made by the following named inventors:

a. Name: John Newlin
Residence (City): Sunnyvale, California
Post Office Address: 718 Old San Francisco Rd., #326
Sunnyvale, CA 94086

Country of Citizenship: United States of America

b. Name: Albert Wang
Residence (City): Los Altos, California
Post Office Address: 871 Gardenia Way
Los Altos, CA 94024

Country of Citizenship: United States of America

c. Name: Christopher M. Songer
Residence (City): San Jose, California
Post Office Address: 1702H Meridian Ave.
San Jose, CA 95125

Country of Citizenship: United States of America

1100 New York Avenue, N.W.
Ninth Floor, East Tower
Washington, D.C. 20005-3918
Tel: (650) 233-4552
Atty/Sec: RSJ/jw
Fax: (650) 233-4545

PILLSBURY MADISON & SUTRO LLP

By: Roger S. Joyner, Reg. No. 36,176

Express Mail Label:
Date of Deposit:

EL 618986055 US
October 3, 2000

I certify that this paper and listed enclosures are being deposited with the U.S. Post Office "Express Mail Post Office to Addressee" under 35 CFR 1.10 on the above date, addressed to Asst. Commissioner for Patents, Box Patent Application, Washington, D.C. 20231

Jeanette Walker
Jeanette Walker

DEBUGGING APPARATUS AND METHOD FOR SYSTEMS OF CONFIGURABLE PROCESSORS

CROSS-REFERENCE TO RELATED APPLICATIONS

5 This application is related to United States Patent Application Nos. 09/246,047 to Killian et al.; 09/323,161 to Wilson et al.; 09/192,395 to Killian et al.; 09/322,735 to Killian et al.; 09/506,502 to Wang et al.; and 09/506,433 to Songer et al., all of which are hereby incorporated by reference.

10 BACKGROUND OF THE INVENTION

1. Field of the Invention

15 The present invention is directed to software development tools for processors and collections of processors that have easily configurable features. In particular, the present invention is directed to software development tools that observe and display the state of a configurable processor or collections of such processors.

2. Background of the Related Art

20 This invention relates to configurable processors as described in the above-referenced Wang et al. application. To summarize, most processor architectures are fixed. The instructions that a processor can execute are fixed, or limited to a small set of variations, and the information, or state, that they maintain is also fixed or limited to a small set of variations. There are new processors, however, that allow the user to change the architecture of the processor including, but not limited to, addition of instructions and of state. This invention relates to the

field of configurable processors. Specifically, the configurability of a processor poses certain problems in software development tools.

This invention also relates to the field of embedded software development. Many computer systems have a CRT for displaying information, a processor, some memory, some sort of fixed storage, a keyboard and other peripherals. These computer systems are equipped to easily display information to the user. Further, they tend to be able to store, run and display many different programs simultaneously. For example, users of Windows computers can run and view many different software programs at the same time.

Other computing systems often have far less hardware associated with them and are far less capable. They often have neither a monitor, a keyboard nor a disk drive. In the simplest case, they may even be completely implemented on a single piece of silicon. Often these systems are capable of storing and running only a single program and have little capability for the visual display of information. Such computing systems are commonly called "embedded systems" because they are embedded in some other system.

Software rarely works properly in its first revision and software developers use a variety of software tools to be able to observe and diagnose software problems. One of the most essential of these tools is a "debugger." A software program is written in program code, hereinafter referred to as code. Code describes both state and operations to be performed on that state. Generally, the user of a software program cannot view the complete state of that program. This is especially true of embedded software programs. Debuggers allow the software developer to view that state as well as how the state is changed by the operations that are performed on it.

Certainly, debuggers are not the only software development tools that require access to the state of the processor. Though they serve as good example of this class of tool,

certain other types of software tools also face the same problems that are faced by debuggers.

For example, monitors provide direct visibility to the state of the processor without

understanding of the programmatic context and so require this information. Silicon co-

verification tools also need this information. Real-time trace capture solutions can need this

5 information. All of these, like debuggers, are generally trying to display state of the processor to

the user and some of them are unique – and so do not fall into some particular class of software

application. Take as a concrete example the software development environment that is provided

by WindRiver in the Tornado 2.0 product. This environment is composed of a variety of

different tools. One of these tools is a debugger called CrossWind. But Tornado 2.0 also

10 contains a tool called the "Browser" that allows visibility to the state of the processor. Another

tool provided with Tornado 2.0 is called "WindShell" and, again, allows the user to access and

view the state of the processor.

In the remainder of this document, the term "debugger" is used to denote any

software tool that requires knowledge of the state of the processor to correctly perform its

15 function, whether it performs a debugging function as such is commonly known in the art or not.

The processor and surrounding hardware maintain the state of the program. Some

of the state is held in registers that are in the processor, but some of the state can be held in

memory or even disk or device registers. So, viewing the state of the program requires viewing

the state of the processor and the surrounding hardware. Consequently, tools of this type must be

20 able to access the state of the processor.

The state of software running on a traditional computer system can be viewed by

debugging programs running on the same traditional computer system. This is because the

computer system is capable of executing and displaying the output of several programs

simultaneously. Embedded systems are usually not capable of running and displaying a debugging program. As a result, embedded systems are usually debugged remotely. A remotely debugged system does not itself execute the debugger software but instead provides some mechanism for debugging software running on another computer system to query information about the state of the system. There are a variety of mechanisms that are used to access the state of the processor in remotely debugged systems, but they generally fall into one of three classes.

The first class of mechanisms involves the software tool, a communication channel to the processor and a program called a monitor that runs on the processor.

Conceptually, processors generally execute instructions that are referenced by address. The address of the currently executing instruction is called the program counter and the processor executes a series of instructions (a program) by executing instructions from different addresses. The processor executes the monitor by the program counter pointing to one of the addresses containing the monitor. The processor then uses its standard instruction execution mechanism to execute the instructions in the monitor. The monitor retrieves the state of the processor and sends that state to the software tool through the communication channel.

The need to execute the monitor through the standard mechanism of instruction fetch limits the information that can be returned to the software tool. Because the monitor itself uses much of the processor to execute, there are times where the processor is in a condition that the monitor cannot be executed, and therefore, cannot make the state visible to a software tool.

To understand this, consider the debug monitor used by systems running the WindRiver VxWorks operating system. This operating system uses a debug monitor that runs as a task under the operating system. Because the debug monitor runs as a task, it can only report information to the debugging software tool when a task is allowed to run. The operating system

does not allow tasks to run when tasks of a higher level are running or when interrupts are being handled. The debug monitor task is also not able to run until a certain point in the initialization process. Consequently debuggers using this debug monitor cannot debug code that is run during the early phases of initialization, that is handling interrupts or that is running at a higher priority than the debug monitor task.

These limitations of monitor programs substantially affect the ability of the debugging tool to aid in diagnosing errors. The other mechanisms to view the state of the processor help to address this limitation.

Before considering the next mechanism, consider the implications of configurable processors (such as those described in the above-referenced applications) for monitor programs. Access to the state of the processor can become problematic in the case of a configurable processor because a monitor, which is responsible for sending state of the processor to the communication channel, must itself know about the state of processor. This problem has been solved in the past by generation of a new monitor for each different processor configuration by the processor generator. The initial release of the Xtensa processor dealt with generation of the monitor in this way. Further, multiple versions of the WindRiver monitor were created for the analogous problem of Intel 486 processors with and without floating point instructions.

This too has problems. Monitor programs are generally stored in ROM or EPROM devices. Programming these devices with a new version of a program generally consumes either additional time or money or both. Further, monitor programs must exist in the same address space as the program to be debugged. A dynamically generated monitor program will vary in its size. Such variation requires changes in additional software for the system to be able to optimally use the available memory space. The present invention addresses this problem.

Before moving to the next class of remote-debugging mechanism, there is one more point to make for remote-debugging using instruction insertion mechanisms. Instruction-insertion remote debugging solutions may use more than one additional computing system. Because instruction-insertion mechanisms require specialized hardware interfaces to access their functionality, that interface is sometimes handled by a different computing system. In this situation, a computing system running the software debugging tool has a communication channel to a server software program that has a different communication channel to the instruction-insertion mechanism. The role of this server program is to translate requests from the software debugging tool into commands to the instruction-insertion mechanism. The server software faces some of the problems faced by the monitor program in this system. The server software must know about the new processor state and new instructions to access that state. One solution to this problem is to rebuild the server software for each processor. But again, the server software faces some of the same issues of convenience and efficiency. The present invention offers a solution to this problem.

The final class of mechanism for a remote debugger to access the state of the processor is through use of direct state scanning. Some hardware options allow the state of the processor to be read directly out of the processor without execution of any instructions (inserted or otherwise).

Fabrication of application-specific chips makes the use of configurable processors possible. Often these application-specific chips will have more than one processor on them. The system designer determines the number of processors in this collection and the arrangement of those processors and their relationship to non-processor elements. Remote software debugging

solutions must work in this type of environment. This invention solves problems for debugging software that arise from this environment.

SUMMARY OF THE INVENTION

5 The present invention has been made with the above problems of the prior art in mind, and it is an object of the present invention to provide a method that allows a single monitor program to support a variety of processor configurations that have different processor state and instructions.

10 It is a further object of the present invention to provide a method that allows a single instruction-insertion server program to support a variety of processor configurations that have different processor state and instructions.

15 It is a further object of the present invention provide a method that allows a remote debugging solution to work for collections of processors, regardless of the number or arrangement of the elements in the processing system.

20 It is another object of the present invention to provide a system and method which permits the state of configurable embedded processors to be easily read, manipulated and debugged by a debugging system.

 It is yet a further object of the present invention to provide a system and method which permits single tasking processor state to be easily read, manipulated and debugged by a debugging system.

 It is still another object of the present invention to provide a system and method which permits reading and manipulation of processor state and other debugging functions in a processor which has a configurable architecture and a variable state structure.

It is a further object of the present invention to provide a system and method which minimizes target processor debugging instruction execution in a debugging system.

It is another object of the present invention to provide a system and method which can perform debugging operations on a configurable processor or system containing configurable processors in spite of active interrupts, high priority level interrupt routines, and initialization programs.

It is still a further object of the present invention to provide a system and method for debugging a processor which can accommodate a wide variety of types of processor state.

It is yet another object of the present invention to provide software for debugging a configurable processor system which does not need to be reconfigured or recompiled for different configurations of the processor.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a diagram of GDB/XMON system topology according to a preferred embodiment of the present invention;

Figure 2 is a diagram of GDB/XOCD system topology according to a preferred embodiment of the present invention; and

Figure 3 is a diagram of a JTAG interface system topology according to a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF
PRESENTLY PREFERRED EXEMPLARY EMBODIMENTS

In perhaps its broadest aspect, there are two facets to the present invention. First, a preferred embodiment of the present invention provides a software system that includes debugging software capable of transmitting instruction sequences for processor state access to either monitor or instruction-insertion server software. Second, a preferred embodiment of the present invention provides a state access mechanism that at run-time reads and understands the structure of a multi-element processing system to allow debugging software to individually identify and access the state of the elements.

In a preferred embodiment of the present invention, a version of the GDB debugger from the GNU family of open source software development tools is able to remotely debug configurable Xtensa processors either individually or in a system of processors. This version of GDB is able to do so without changes to either the monitor program or the instruction-insertion server. It does this using information created with the Xtensa configurable processor development system. The processors themselves are preferably created with the Xtensa configurable processor development system of Tensilica, Incorporated of Santa Clara, California which is generally described in the aforementioned patent applications.

Hereinafter, “GDB” will denote a version of GDB modified according to the teachings of the present invention.

Single Core Debugging

Debugging a single configurable core with GDB explores the first facet of this invention. As described in the above-cited applications, the processor generator creates the

configurable core with input from the user. This input from the user includes a definition of additional processor state to be included in the processor as well as the instructions that move that state into memory.

The additional state and register information is described in the TIE language as documented in the Tensilica TIE Reference Manual, incorporated by reference. As a part of the TIE compilation process, the TIE compiler parses the language and identifies the `loadi` and `storei` directives included in the information. These directives give the TIE compiler the information necessary to know how to save and restore a given piece of additional processor state.

The load or store sequence for additional state can itself require the use of other pieces of additional state. As an example, consider a case where the user has declared two register files A and B. The store instructions for register file A use the standard Xtensa address registers for the memory address to which to store. However, the store instructions for register file B use the values of register file A for the memory address to which to store. Therefore, the store of a register from register file B first requires the store of a register from register file A.

The TIE compiler generates this dependency information and encodes it in the form of several dynamically loadable libraries called `libcc`'s. A `libcc` is a library that can be loaded at program execution time by other software libraries. It includes information that is processor-specific. In particular, `libcc` contains information regarding the save and restore information and sequences for given registers as well as the name and size information. It is worth noting that the syntax and semantics of the TIE language require that the register file save dependencies not be circular. The TIE compiler will generate an error in the presence of `loadi/storei` directives that create a circular dependency graph.

GDB calls a library called `libdb` for information about a given processor configuration. `Libdb` is C library that has functions that provide information about the state of the processor. This library is compiled into a library that is used by an application that wants to access the information. Consider the following example of GDB's use of `libdb`. First, we will

5 look at the initialization of `libdb` as GDB gets the register information for a particular processor:

```

static void
init_libdb()
{
10     int index = 0;

    g_debug_object = xtensa_init_debug_object(
xtensa_default_params );

15     NUM_LIBDB_REGS = xtensa_get_register_count(
g_debug_object );

    for (index = 0; index < NUM_LIBDB_REGS; ++index)
    {
20         libdb_reg_map[index] = xtensa_get_register_info(
g_debug_object, index);
    }
}

```

25 Note that the debugger does not have to know anything *a priori* about the registers that are in the machine. Instead, it is getting all of the register information from `libdb`. (In the above code sequence, the following calls were `libdb` calls: `xtensa_init_debug_object`, `xtensa_get_regster_count` and `xtensa_get_reg_info`.)

30 Once the initialization has been performed, information about the processor can be acquired from `libdb`. For example, the following code snippet looks up register information by register name:

```
reg_map[n].libdb_number = xtensa_find_register_by_name(  
g_debug_object, reg_map[n].name );
```

This library, in turn, loads and calls the `libcc`'s produced by the TIE compiler.

- 5 Libdb is able, at the request of GDB, to generate instruction sequences to save and restore any piece of state that has been added to the processor.

Before dealing with another example, let us discuss at a high level what happens.

If each of the register files is considered a node and each dependency considered a directed edge from one node to another, then the dependencies of the save and restore instructions form a

- 10 directed acyclic graph (DAG) (recall that the TIE compiler has guaranteed that the dependencies are acyclic). Libdb forms the instruction sequence to save and restore each piece of state by depth-first traversal of the DAG (starting at state to be saved and/or restored). In this depth-first traversal, each child node is visited and then the instructions for that node are generated. A given set of instructions can have multiple representations. The assembly source code for the
15 instructions is one form of those instructions, but the processor cannot execute this form. The processor can only execute the machine form of these instructions. So, `libdb` generates the machine form and makes that form available to GDB.

Consider the following example. A particular TIE coprocessor has the following register file and compiler type definitions (an explanation of the TIE language may be found in,

- 20 e.g., the above-referenced Tensilica TIE Reference Manual):

```
regfile vec 160 16 v  
regfile align 112 4 u
```

- 25 ctype vec8x20 160 128 vec
ctype align 128 128 align

These statements declare two register files `vec` and `align`. The register file `vec` is assigned the ctype `vec8x20` and `align` is assigned the ctype `align`. These types have the following associated `loadi` and `storei` directives.

```

5  proto align_loadi {out align u, in align* s, in immediate o} {vec8x20 t} {
    LV16.I      t, s, o;
    WALIGN      u, t;
  }
  proto align_storei {in align u, in align* s, in immediate o} {vec8x20 t} {
10    RALIGN      t, u;
    SV16.I      t, s, o;
  }
  proto vec8x20_loadi {out vec8x20 t, in vec8x20* p, in immediate o} {} {
    LV32.I      t, p, o;
    LVH.I       t, p, o+16;
15  }
  proto vec8x20_storei {in vec8x20 t, in vec8x20* p, in immediate o} {} {
    SVL.I       t, p, o;
    SVH.I       t, p, o+16;
20  }

```

The `align_loadi` proto directive defines the instruction stream necessary to load a value into the alignment register file. Note that the `LV16.I` and `WALIGN` instructions require the use of a `vec8x20` typed register (a `vec` register). In the same way, the `align_storei` directive defines the instruction stream necessary to save a value from the `align` register file. In the same way, the `storei` directive uses instructions that require the use of a `vec8x20` register. Both saving and restoring an `align` value requires the use of a free `vec8x20` register.

Note that if each type is considered a node and each dependency is considered an edge then this example forms a graph with the set of nodes `{align, vec8x20}` and the set of edges `{(align, vec8x20)}`. The TIE compiler parses the source above and forms a dependency graph from the declarations. It then checks this graph for circular dependencies through standard graph algorithms known in the art such as the kind in Introduction to Algorithms (Cormen, Leiserson and Rivest). The TIE compiler then generates C code that

represents these non-circular proto directives. The TIE compiler further encodes the save and restore instructions themselves into this C code. This C code is compiled into a `libcc` file.

`Libddb` loads the `libcc` and accesses the encoded dependency information. To generate the save information for an `align` register, `libddb` looks at the `proto` information for the `align` register that was put into the `libcc` by the TIE compiler. This information gives both the instructions and the dependencies. Seeing that there is a dependency on the `vec8x20` type, `libddb` first generates a set of instructions (from the instruction information encoded in the `libcc` file by the TIE compiler) to save a `vec8x20` register. It then generates a set of instructions, using this saved `vec8x20` register to save the alignment register. Finally it generates a set of instructions to restore the `vec8x20` register from the save locations. This is because the saving of the state is not guaranteed to leave the values of the intermediate registers intact. As a consequence, the intermediate values must be restored so that the operation does not disturb the state of the processor.

To summarize: GDB gets the save and restore sequence from `libddb` while `libddb` generates the save and restore sequence by using information that has been produced by the TIE compiler. The TIE compiler is a part of the processor generator. Of course, the whole reason that GDB is requesting information about a particular register's value is that source level debuggers are showing the state of the processor to the user. Compilers (and in this case, the `xt-gcc` Xtensa-C/C++ compiler used to generate the configurable processor architecture) not only generate machine code from source code, but also generate extra information to tell tools such as debuggers where a given piece of state is stored at any given time. When the user wishes to view a particular piece of state in code, the debugger looks at this extra information to determine where that information is stored. Sometimes the information is in memory and the debugger

reads the state from memory. Sometimes, however, the state is in registers and the debugger must read the state directly from the register. It is also worth noting that users can directly request the register information and this can be of direct use to the user.

Users of GDB can remotely debug an Xtensa processor with the monitor XMON.

- 5 XMON is the Xtensa monitor program and runs on Xtensa processors. The details of the XMON monitor involve the intricacies of the Xtensa architecture so please refer to the Xtensa Instruction Set Architecture Reference Manual, incorporated by reference, for details. Source code for a version of XMON is included in Appendix A.

10 XMON is a software debugging monitor that uses the processor to access the state of the processor for GDB. XMON communicates with GDB over a serial link. XMON remains the same for versions of the processor that have different state and instructions to retrieve that state (see FIG. 1). XMON is kept the same to lower the amount of work required to deliver a running system. Building a new monitor requires burning new ROMs to install that monitor and there is an efficiency to having a single monitor that can service multiple processor configurations. Though changes to other aspects of the processor can require a new XMON to be built, a given XMON works properly for all possible additional state and state retrieving instructions. Changes that require a new XMON to be built are changes that affect the communication mechanisms that XMON uses or changes that affect the actual format of the core instructions that the processor can execute. For example, XMON depends on the serial port to communicate with the host debugger. That serial port is connected to an interrupt on the processor and that interrupt has a particular interrupt level in the processor. Both the interrupt and the interrupt level are configurable by the user. If the user changes either of these parameters for the serial port interrupt, a new XMON must be built. Xtensa processors can also be

15

20

configured to execute either in a big-endian or a little-endian format. Depending on the configuration, the encoding of the instructions changes. As a consequence, a new XMON must be built when the processor endianness is changed.

GDB retrieves new state of the processor from XMON by getting the save

- 5 sequence for that state from libdb. The communication protocol for communications between GDB and XMON is shown in the tables below.

Command	Description	Response
?	Returns the most recent signal received from the target which can be one of: GDB_SIGINT GDB_SIGQUIT GDB_SIGILL GDB_SIGTRAP GDB_SIGIOT GDB_SIGBUS GDB_SIGFPE GDB_SIGSEGV GDB_SIGALRM GDB_SIGTERM	SXX
pRRR...R	Read a single register whose number is RRR...R. The way registers are numbered is implementation and configuration dependent.	XXX...X
PRRR...R=XXX...X	Write the value XXX...X into a single register whose number is RRR...R.	OK E<message>
mAAA...A, LLLL;	Read the LLLL bytes from address AAA...A on the target. Return the contents as a hex string.	XXX...X
MAAA...A, LLLL:XXX...X	Write LLLL bytes at address AAA...A. The value to be written is the hex string XXX...X.	OK E<message>
C	Continue execution at the current pc. Respond with	S<NN> E<message>

	the reason code <NN> when program stops.	
S	Step one instruction, and return the reason code <NN> (in unix lingo, the signal value) that caused it to stop. If the step succeeded NN=05.	S<NN> E<message>
K	Not implemented	
Xqn	This command exists for two reasons. First, xt-gdb issues this command to determine what the kind of target it has connected to. XMON responds with the string "XMON" identifying its version to xt-gdb. This version number is particularly important since xt-gdb will adjust its register numbering to match the version of XMON.	XMON1.5 XMON2.0 XOCD ISS ISS3
XqpXXXXXXXX	Queries the target to see if the target "knows" how to fetch the specified register	Y N
XqpXXXXXXXX	Queries the target to see if it knows how to set the specified register	Y N
XexeXX:YY:ZZ	Executes an arbitrary opcode on the target. The OCD daemon does byteswapping for bigendian targets.	NULL string
Xsbe[0 1]	Tells the OCD daemon that the target is big endian, or little endian.	Null string
XsisXX	Sets the cache line size. This is used so that when the OCD daemon is accessing memory, it can do cache flush instructions on the appropriate byte boundaries, or skip the	Null String

	cache flush instructions in the case of not having an instruction cache.	
Xcs	Toggles stack spilling on and off. By default the xocd daemon will spill the AR register file to the stack. This makes doing stack traces quick and easy, but can make some forms of debugging not work correctly. This toggles the stack spill policy. Use this with extreme caution.	OK
XwgrXX:name:YYYYYYYY	Writes a “generic” tap’s, tap register, where XXX is the device number on the jtag chain, name is the name of the register as specified in the topology.ini file, and YYYY is the value being written.	OK E:Error Writing Register
XqgrXXXX:name	Reads a “generic” tap’s tap register. XX is the device number on the jtag chain, and name is the name of the tap’s register as specified in the topology.ini file.	OK E:Error Reading Register

TABLE I

An extended protocol shown in TABLE II offers more usability. For instance, the extended protocol will support breakpoints in ROM while the base protocol cannot, as all breakpoints are

5 implemented by writing a break instruction into memory.

Command	Description	Response
g	Request that target to return the contents of all known registers. The reply is a hex string; each pair of digits represents a byte; and the	XXXXXX XXXXX

Sending packet: \$p80000e0#fd...Ack
Packet received: 0000000e
Sending packet: \$P80000e0=0000000f#d0...Ack
Packet received: OK

5

3. This saves the memory (4 bytes in this case) where the TIE register will be spilled.

Sending packet: \$m4005b150,4#8e...Ack
Packet received: 00000000

10

4. This saves the a4 register, which will be used as the addressing register. Then a4 is set

to contain the address where the TIE register will be spilled.

Sending packet: \$p4000004#c8...Ack
Packet received: 40010088
Sending packet: \$P4000004=4005b150#a6...Ack
Packet received: OK

15

5. This sends the instruction to be executed, in this case, we use a4 for address, and are
spilling register i321 (a user defined TIE register).

Sending packet: \$Xexe:31:41:00#71...Ack
Packet received:

20

6. Now we read the memory at the location where the register was spilled. This should
be the value of the register, and in this case the register contained the value zero which is what
we read back.

25

Sending packet: \$m4005b150,4#8e...Ack
Packet received: 00000000

7. Now we restore a4, and the memory where we spilled the TIE register.

Sending packet: \$P4000004=40010088#7a...Ack
Packet received: OK
Sending packet: \$M4005b150,4:00000000#28...Ack
Packet received: OK

30

8. Now we restore CPENABLE.

35

Sending packet: \$P80000e0=0000000e#cf...Ac
Packet received: OK

Users of GDB can also remotely debug an Xtensa processor with the processor

feature OCD. OCD is an instruction-insertion feature described in the Tensilica On-Chip Debug

40

Mode User's Guide, incorporated by reference, and XOCD is a server that uses the OCD feature to access processor state. XOCD communicates with GDB across a TCP network connection and communicates with the hardware instruction-insertion mechanism through a special piece of hardware called "the wiggler." The wiggler connects to the parallel port of a PC and converts those signals to electrical signals appropriate for a JTAG connection (see FIG.2). XOCD does not have to be rebuilt for different processor configurations. A single version of the XOCD software will service all processor configurations.

As with XMON, GDB retrieves new state of the processor from XOCD by getting the save sequence for that state from `libdb`. There is a communication protocol between GDB and XOCD. This protocol is the same protocol as for XMON (see TABLE I above). One of the messages in this protocol tells XMON that GDB is going to send a series of instruction over the serial port to the XOCD server and that the XOCD server should use the instruction insertion feature of the processor to execute those instructions.

Consider the following example of an exchange between GDB and the XOCD server:

1. Query the target to see if it can access the specified register. The target replies with "n" meaning no.

Sending packet: `$Xqp10020001#bd...Ack`
Packet received: n

2. This sequences sets the CPENABLE so that the processor has access to all the coprocessors.

Sending packet: `$p80000e0#fd...Ack`
Packet received: 0000000e
Sending packet: `$P80000e0=0000000f#d0...Ack`
Packet received: OK

3. This saves the memory (4 bytes in this case) where the TIE register will be spilled.

Sending packet: \$m4005b150,4#8e...Ack
Packet received: 00000000

4. This saves the a4 register, which will be used as the addressing register. Then a4 is set

5 to contain the address where the TIE register will be spilled.

Sending packet: \$p4000004#c8...Ack
Packet received: 40010088
Sending packet: \$P4000004=4005b150#a6...Ack
Packet received: OK

5. This sends the instruction to be executed, in this case, we use a4 for address, and are
spilling register i321 (a user defined TIE register)

15 Sending packet: \$Xexe:31:41:00#71...Ack
Packet received:

6. Now we read the memory at the location where the register was spilled. This should
be the value of the register, and in this case the register contained the value zero which is what
we read back.

Sending packet: \$m4005b150,4#8e...Ack
Packet received: 00000000

7. Now we restore a4, and the memory where we spilled the TIE register.

Sending packet: \$P4000004=40010088#7a...Ack
Packet received: OK
Sending packet: \$M4005b150,4:00000000#28...Ack
Packet received: OK

8. Now we restore CPENABLE

Sending packet: \$P80000e0=0000000e#cf...Ack
Packet received: OK

Those instructions save the requested state of the processor into memory and then
bring that data out through the scan chain (see the above-referenced On-Chip Debug Mode
User's Guide Tensilica publication for a detailed explanation of the instruction-insertion
mechanism for Xtensa processor cores).

Again, because GDB is transmitting the run-time generated instructions to XOCD across the network, XOCD does not have to be modified for each new processor configuration but works properly for each one.

It is also worth noting that GDB itself does not have to be modified for each configuration. Unlike the system described in the above-cited applications, using this system, one GDB can service all processor configurations by having `libddb` load the appropriate libraries (`libccc's` -- dynamically linked libraries) that were created by the processor generator. In the previous system, a custom GDB was generated for each processor configuration. This had the drawbacks that the time to build the processor was increased and the amount of space required to store the tool chain was increased.

JTAG Overview

The JTAG specification (available from IEEE) specifies both an electrical and architectural interface. The intent of this interface is to give visibility into silicon systems without requiring lots of additional hardware resource. In particular, the JTAG interface is designed to use minimal silicon area and pins. Debugging hardware that uses JTAG can be considered to have a series of controllers that are connected together in a particular order.

The data and control portion of the JTAG interface is a serial bit stream. Both instructions to be performed and the results of those instructions are transmitted over this serial bitstream. As different devices are connected to the same JTAG interface, the total length of the bit stream becomes longer. (See FIG. 3). Each of the objects on the scan chain is a set of logic called a TAP controller and the TAP controller accepts instructions from the serial interface.

Multiple Core Debugging

For a variety of reasons, monitors do not generally provide effective debugging solutions for multiple core debugging. One reason is addressability. The processor needs to know how to select and observe each core individually. However, monitor programs generally use communication channels that are difficult to aggregate and/or split. Take XMON as an example. XMON uses a serial port to communicate with GDB. Serial is a very simple protocol, but it is not designed to have more than two objects on the connection. One obvious solution is to have one serial port per processor, but in the case where many processor cores are implemented on a single piece of silicon, all of those serial ports either have to be aggregated, or each has to come out to pins on the package. The number of pins available on a package is limited and solutions that use fewer pins are superior.

WindRiver Systems' Tornado2.0 environment uses TCP for communication between the monitor and the debugging software. While network connections are easy to aggregate and split, the amount of hardware required to implement this solution is reasonably high and is not efficient on the scale that would be necessary to make them viable for multiple cores on a single piece of silicon.

The IEEE JTAG specification provides an interface that is easy to aggregate, easy to split and reasonable in size. Instruction-insertion mechanisms such as OCD on the Xtensa core tend to use the JTAG interface. In the preferred embodiment, the processor cores are connected together serially in a JTAG chain. The XOCD server is connected to the end point of this chain.

The XOCD server (described in greater detail in the aforementioned Tensilica On-Chip Debug Mode User's Guide) controls the instruction insertion mechanism by shifting

bits into the JTAG scan chain one bit at a time. Because a bit will go from one bit in a register to another bit in a register, and finally to the first bit in the next register, the XOCD server has to know where the processor is in the chain to be able to address a given processor. The simplest way to do this would be to build a custom XOCD server for each system topology. But this solution has a series of drawbacks. A different XOCD server would have to be built and installed for each processor because the XOCD servers would be different for each processor or system of processors.

In the preferred embodiment, a generic XOCD server reads a file that specifies the topology of the system of configurable processors. This file includes information about position in the scan chain for each processor as shown in TABLE III below.

Section	Key	Description
[main]	Number_of_xtensas	Number of Xtensa processors in the scan chain.
	Number_of_generic	Numer of instances of the generic TAP controller.
	Number_of_other	Number of other (bypassed) TAP controllers
[generic_description]	IR_Width	Instruction register length in bits.
	Bypass	Bypass instruction
	Bypass_length	Bypass data register length
	Number_of_regs	Number of accessible data registers
[generic_reg_X]	Gdb_name	Name given to the data register for the debugger access.
	Width	Data register width in bits.
	Read_Instruction	Instruction that selects the data register
	Write_Instruction	Instruction that selects the data register (if writable)
[xtensa_X]	Position	Position of the Xtensa on the TAP chain.
[genic_X]	Position	Position of the generic TAP instance on the chain.
[other_X]	Position	Position of the bypassed TAP controller on the chain.
	IR_Width	Instruction register length in bits.

Bypass	Bypass instruction.
Bypass_length	Bypass data register length.

TABLE III

To best understand the use of the topology file, let us consider several files. First a file that has only a single processor on the scan chain.

```
[main]
number_of_xtensas = 1
number_of_generic = 0
number_of_other   = 0
```

```
[xtensa0]
position=0
```

Clearly there is not much to say about this example. With only one Xtensa, and no other JTAG TAP interfaces on the chain, the Xtensa must be at position 0.

In the next example, there are two Xtensa processors.

```
[main]
number_of_xtensas = 2
number_of_generic = 0
number_of_other   = 0
```

```
[xtensa0]
position=1
```

```
[xtensa1]
position=0
```

This example illustrates that the position of the processor does not have to match the numbering of the processor. When GDB connects to the XOCD server, it connects to a particular TCP/IP socket that is listened to by the XOCD server. That socket corresponds to the Xtensa number rather than the position number.

Of course, a scan chain can have more than just processor controllers on it. In the preferred embodiment, the topology file also specifies additional controllers. This specification includes both their impact on the scan chain, their basic accessibility information and their architectural information. The XOCD server takes this information and allows multiple instances of one particular TAP controller architecture to be addressed by the debugger. This allows the user to access state that is not associated with the processor core, but is instead associated with the surrounding system.

In the next example, the two Xtensa processors are joined by an additional TAP controller that is not an Xtensa TAP controller.

```
[main]
number_of_xtensas = 2
number_of_generic = 1
number_of_other   = 0

[xtensa0]
position=1

[xtensa1]
position=0

[generic_description]
IR_Width = 5
bypass   = 0x1f
bypass_length = 1
number_regs = 2

[generic_reg_0]
gdb_name=MOTOR_SPIN
Width=2
Read_Instruction = 0x17
Write_Instruction = 0x18

[generic_reg_1]
gdb_name=SPIN_RATE
Width=32
Read_Instruction = 0x19
Write_Instruction = 0x20
```

```
[generic0]
position=2
```

5 The additional TAP controller is a “generic” TAP controller that is defined to have two registers, the MOTOR_SPIN register and the SPIN_RATE register. The definitions of these registers include all the information that is necessary for the XOCD server to access these registers and return the information to the user.

10 In the final example, the two Xtensa processors are joined by an additional TAP controller that is not accessible to the XOCD server. The user has only described enough information about the TAP controller for the XOCD server to be able to use the scan chain. The particulars of the TAP itself are not described. This contrasts with the description of the above TAP controller.

```
15       [main]
         number_of_xtensas = 2
         number_of_generic = 0
         number_of_other   = 1
```

```
20       [xtensa0]
         position=1
```

```
25       [xtensa1]
         position=0
```

```
30       [other0]
         position=2
         IR_Width=5
         bypass=0x1f
         bypass_length=1
```


WHAT IS CLAIMED IS:

1. A method of accessing state from a configurable processor, the method comprising:
transmitting, using a software application, a state-accessing instruction stream to an agent in
the configurable processor, the software agent being capable of interpreting that stream; and
causing, using the state-accessing instruction stream, the interpreting agent to return the
state of the processor to the software application.
2. A method as in claim 1 where the interpreting agent is a monitor program.
3. A method as in claim 1 where the interpreting agent is an instruction insertion server.
4. A method as in claim 1 where the interpreting agent is an architectural simulator.
5. A method as in claim 1, further comprising:
reading, using the software application, information describing the configurable processor's
state architecture; and
generating, using the software application, the instruction stream based on the information.
6. A method as in claim 5 wherein the interpreting agent is a monitor program.
7. A method as in claim 5 wherein the interpreting agent is an instruction insertion server.
8. A method as in claim 5 wherein the interpreting agent is an architectural simulator.

9. A computer-readable storage medium storing therein a software program capable of generating a processor from a user description of that processor that also generates information necessary to describe save and restore instructions for state of the processor.

10. A computer-readable storage medium storing therein a software library usable for reading a description of save and restore information and then generating saving and restoring instruction streams therefrom.

11. A medium as in claim 10 wherein the software library also can deal with interdependencies in state to generate a complete and correct save and restore sequence.

12. An instruction-insertion server that takes system topology information from a computer-readable file to determine where elements are in a system described by the file.

13. A system for accessing state from a configurable processor, the system comprising:
a software application which transmits a state-accessing instruction stream;
an interpreting agent in the configurable processor which
receives the instruction,
interprets the stream to access state of the configurable processor, and
returns the accessed state of the configurable processor to the software application.

14. A system as in claim 13 where the interpreting agent is a monitor program.

[illegible][illegible]

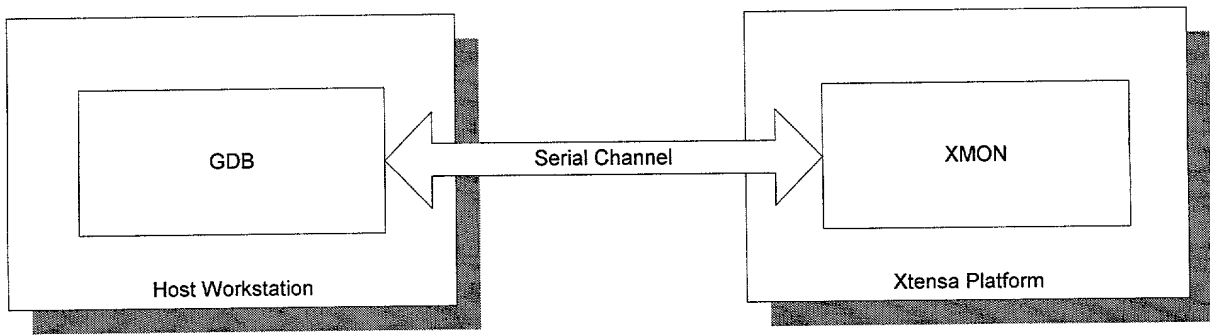


FIGURE 1

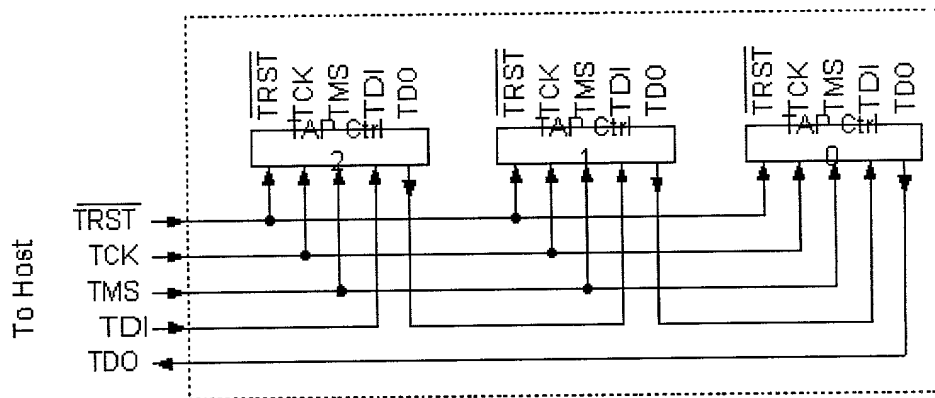


FIGURE 3

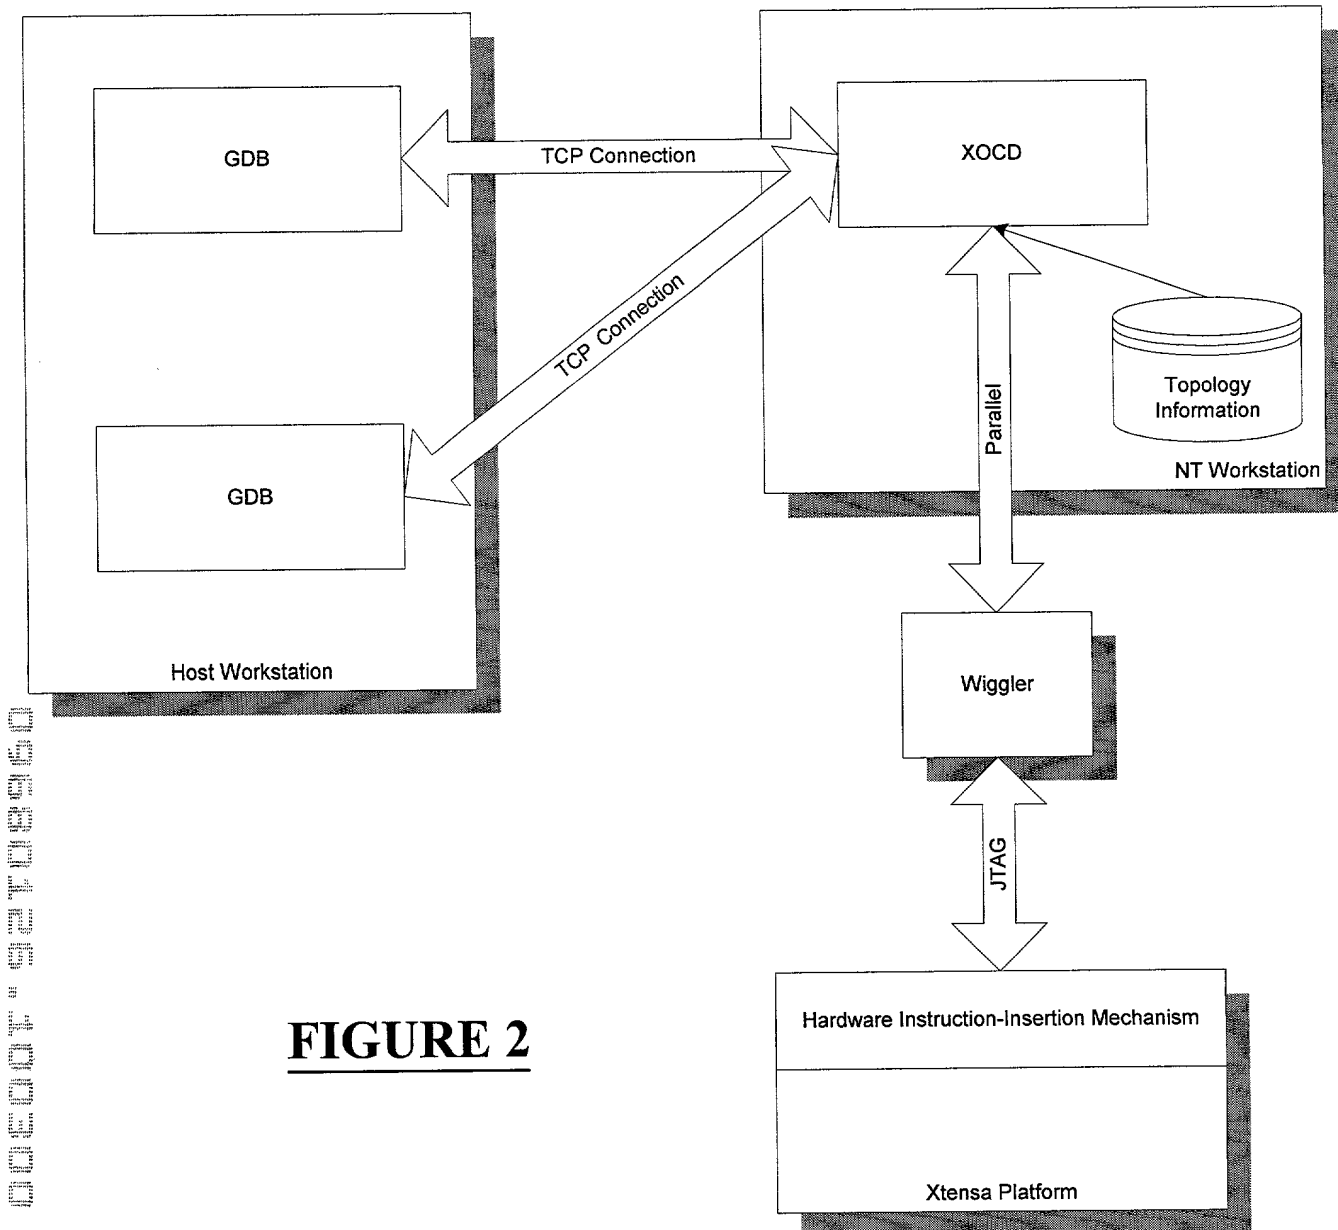


FIGURE 2

Appendix A, example XMON source Code

XMON primarily consists of two parts. The first part handles the Debug exception of the Xtensa processor. This is implemented in the file `DebugExceptionVectorHandler-mon.S`. The second part is implemented in `xtensa-mon.c` and handles the higher level protocol with the debugger.

I. DebugExceptionVectorHandler-mon.S.

```
// Exports
.global _DebugExceptionFromVector
.global _ar_registers
.global _sr_registers
.global _level_0_interrupt
.global _flush_i_cache
.global _xmon_out
.global _xmon_in
.global _xmon_flush
.global _xmon_init

// Imports
//      _handle_exception

#include <machine/specreg.h>
#include "DebugExceptionVectorHandler.h"

#define AR_SAVE_SIZE (4*NUM_AREGS)
#define SR_SAVE_SIZE (4*256)

// Parameters
#define XMON_STACK_SIZE (2048+1024)
```

60183116_1.DOC

/******

The assembler portion of the debug handler begins here. The handler does three major things. First, it saves the processor state. The bulk of the save sequence saves all the address registers. Note that we don't try to save the registers into interrupted process' stack because it may have become corrupted and the debugger wants to perturb the processor state as little as possible. Second, the handler sets up the run-time environment for the debugger stuff, which we have written in C. Third, upon return from the stub, we restore the interrupted process' registers, and resume the process. The debugger can force the process to resume at an alternative pc by overwriting the saved value of the appropriate EPC.

In the comments below, we will use "ipwb" to refer to the interrupted process' window base, and "wb" to the current window base.

*****/

```
//      .section .bss
//      .section .text
//      .align 16
//_ar_registers:
//      .space AR_SAVE_SIZE
//_sr_registers:
//      .space SR_SAVE_SIZE
_xmon_stack:
    .space XMON_STACK_SIZE
_xmon_stack_bot:

    .text
    .begin literal
    .align 4
_xmon_stack_ptr:
    .word _xmon_stack_bot-4*16
ar_save_ptr:
    .word _ar_registers
sr_save_area_ptr:
    .word _sr_registers
//ar_save_area_ptr:
//    .word _registers+AR0_OFFSET

    .globl _handle_exception
handler:
    .word _handle_exception

    .align 4
.Laddress_of_save1_table_ptr:
    .word save1_table_ptr

    .align 4
save1_table_ptr:
    .word save1_28 /* ipwb=0 */
    .word save1_24 /* ipwb=1 */
    .word save1_20 /* ipwb=2 */
    .word save1_16 /* ipwb=3 */
    .word save1_12 /* ipwb=4 */
    .word save1_8 /* ipwb=5 */
    .word save1_4 /* ipwb=6 */
    .word save1_0 /* ipwb=7 */

    .align 4
.Laddress_of_save2_table_ptr:
    .word save2_table_ptr
save2_table_ptr:
    .word save2_0 /* ipwb=0 */
    .word save2_4 /* ipwb=1 */
    .word save2_8 /* ipwb=2 */
    .word save2_12 /* ipwb=3 */
    .word save2_16 /* ipwb=4 */
```



```

.word    save2_20      /* ipwb=5 */
.word    save2_24      /* ipwb=6 */
.word    save2_28      /* ipwb=7 */

.align 4
.LDWOE:
.word 0xffffbfff

.Lps:
.word (1<<18)          /* WOE and KM */
.Laddress_of_restore1_table_ptr:
.word    restore1_table_ptr
restore1_table_ptr:
.word    restore1_28    /* ipwb=0 */
.word    restore1_24    /* ipwb=1 */
.word    restore1_20    /* ipwb=2 */
.word    restore1_16    /* ipwb=3 */
.word    restore1_12    /* ipwb=4 */
.word    restore1_8     /* ipwb=5 */
.word    restore1_4     /* ipwb=6 */
.word    restore1_0     /* ipwb=7 */

.align 4
.Laddress_of_restore2_table_ptr:
.word    restore2_table_ptr
restore2_table_ptr:
.word    restore2_0     /* wb=0 */
.word    restore2_4     /* wb=1 */
.word    restore2_8     /* wb=2 */
.word    restore2_12    /* wb=3 */
.word    restore2_16    /* wb=4 */
.word    restore2_20    /* wb=5 */
.word    restore2_24    /* wb=6 */
.word    restore2_28    /* wb=7 */
.end literal

.text
.align 4

```

_DebugExceptionFromVector:

```

/* Save a0,a1,a2 into various places so that we can setup
the save sequence. Notice that we need to take care
that this code works even when a0, a1 contain the
same value. See the NOOP comments.
*/
l32r    a0,ar_save_ptr
s32i    a1,a0,4
s32i    a2,a0,8
l32r    a2,sr_save_area_ptr
rsr     a1,WINDOWSTART
s32i    a1,a2, (WINDOWSTART*4) /* save windowstart */
rsr     a1,WINDOWBASE
s32i    a1,a2,(WINDOWBASE*4) /* save WB */
slli    a1,a1,4              /* multiply by 16, size in bytes of
                             the 4 register window */

add     a1,a0,a1
/* At this point:
a0: address of save area.
a1: &save_area+wb*16 (i.e. save area for current window )
We must ensure that code below works even when a0==a1
*/
rsr     a2,EXCSAVE_0
s32i    a2,a1,0              /* save a0 */
l32i    a2,a0,4
s32i    a2,a1,4              /* save a1; NOOP if a0==a1 */
l32i    a2,a0,8
s32i    a2,a1,8              /* save a2; NOOP if a0==a1 */
s32i    a3,a1,12             /* save a3 */
/* Now save other windows.
We use jump tables to do this.

```

First, we save windows wb+1...n-1 where n == number of windows.
Second, we save windows 0,...,wb-1

```

/*
/* Disable WOE */
l32r    a3, .LDWOE
rsr     a2, PS
and     a2, a2, a3
wsr     a2, PS
rsync

addi    a0,a1,16      /* compute next save area */
rsr     a1,WINDOWBASE
slli    a1,a1,2
l32r    a2,.laddress_of_save1_table_ptr
add     a2,a1,a2
l32i    a2,a2,0
jx      a2
/* The instruction jumps into the 1st part of the save sequence
   with the following notable register contents:
   a0 = ar_save_area_ptr + (ipwb+1)*16; i.e the save area for the
                                   next window.

   wb = ipwb
*/
save1_28: /* ipwb = 0 */
s32i    a4,a0,0
s32i    a5,a0,4
s32i    a6,a0,8
s32i    a7,a0,12
addi    a4,a0,16
rotrw   1
save1_24: /* ipwb = 1 */
s32i    a4,a0,0
s32i    a5,a0,4
s32i    a6,a0,8
s32i    a7,a0,12
addi    a4,a0,16
rotrw   1
save1_20: /* ipwb = 2 */
s32i    a4,a0,0
s32i    a5,a0,4
s32i    a6,a0,8
s32i    a7,a0,12
addi    a4,a0,16
rotrw   1
save1_16: /* ipwb = 3 */
s32i    a4,a0,0
s32i    a5,a0,4
s32i    a6,a0,8
s32i    a7,a0,12
addi    a4,a0,16
rotrw   1
save1_12: /* ipwb = 4 */
s32i    a4,a0,0
s32i    a5,a0,4
s32i    a6,a0,8
s32i    a7,a0,12
addi    a4,a0,16
rotrw   1
save1_8: /* ipwb = 5 */
s32i    a4,a0,0
s32i    a5,a0,4
s32i    a6,a0,8
s32i    a7,a0,12
addi    a4,a0,16
rotrw   1
save1_4: /* ipwb = 6 */
s32i    a4,a0,0
s32i    a5,a0,4
s32i    a6,a0,8
s32i    a7,a0,12
rotrw   1

```



```

#define SAVE(r) \
    rsr    a2, r; \
    s32i    a2, a0, (r*4)

#ifdef ACCLO_OFFSET
    SAVE(ACCLO)
    SAVE(ACCHI)
    SAVE(MR_0)
    SAVE(MR_1)
    SAVE(MR_2)
    SAVE(MR_3)
#endif

#ifdef AV0_OFFSET
    SAVE(AVLO)
    SAVE(AVHI)
    SAVE(BV)
    SAVE(SAV)
#endif

#ifdef BR_OFFSET
    SAVE(BR)
#endif

    SAVE(CACHEATTR)

#ifdef CCOUNT_OFFSET
    SAVE(CCOUNT)
#endif

#ifdef CPENABLE_OFFSET
    SAVE(CPENABLE)
#endif

    SAVE(DEBUGCAUSE)
    SAVE(EPC_1)
    SAVE(EXCSAVE_1)
    SAVE(EXCCAUSE)
    SAVE(ICOUNT)
    SAVE(ICOUNTLEVEL)
    SAVE(INTENABLE)
    SAVE(INTREAD)
    SAVE(LBEG)
    SAVE(LCOUNT)
    SAVE(LEND)
    SAVE(SAR)

/* Disable Interrupts and Icounts */
movi.n a2, 0
wsr    a2, INTENABLE
wsr    a2, ICOUNTLEVEL
wsr    a2, ICOUNT
isync

/* Load new PS: Enable WOE and lower priority.
   We have already turned off interrupts and icount
   from above. */
l32r    a1, .Lps
wsr    a1, PS

movi.n a0, 0
movi.n a2, 1
wsr    a2, WINDOWSTART          /* window start = 1 */
wsr    a0, WINDOWBASE          /* window base = 0 */
rsync

/* Initialize our stack and call handler */
l32r    a1, _xmon_stack_ptr
l32r    a2, handler
callx4 a2

/* Raise interrupt level back up and disable WOE */
rsr    a2, PS
movi.n a3, 0
or     a2, a2, a3
l32r    a3, .LDWOE
and    a2, a2, a3

```

```

        wsr      a2, PS
        rsync

        /* restore sequence */
        l32r     a0, sr_save_area_ptr

#define RESTORE(r)      \
        l32i     a2, a0, (r*4);      \
        wsr      a2, r

#ifdef ACCLO_OFFSET
        RESTORE(ACCLO)
        RESTORE(ACCHI)
        RESTORE(MR_0)
        RESTORE(MR_1)
        RESTORE(MR_2)
        RESTORE(MR_3)
#endif
#ifdef AV0_OFFSET
        RESTORE(AVLO)
        RESTORE(AVHI)
        RESTORE(BV)
        RESTORE(SAV)
#endif
#ifdef BR_OFFSET
        RESTORE(BR)
#endif
        RESTORE(CACHEATTR)
#ifdef CCOUNT_OFFSET
        RESTORE(CCOUNT)
#endif
#ifdef CPENABLE_OFFSET
        RESTORE(CPENABLE)
#endif
        RESTORE(EPC_1)
        RESTORE(EXCSAVE_1)
        RESTORE(EXCCAUSE)
        RESTORE(INTENABLE)
        RESTORE(INTREAD)
        RESTORE(LBEG)
        RESTORE(LCOUNT)
        RESTORE(LEND)
        RESTORE(SAR)

        /* Now restore all the ar's */
        l32i     a2, a0, (WINDOWBASE*4)
        wsr      a2, WINDOWBASE      /* set wb to ipwb */
        rsync
        l32r     a0, ar_save_ptr
        rsr      a2, WINDOWBASE
        slli     a2, a2, 2            /* multiply by 4 */
        l32r     a3, .laddress_of_restore1_table_ptr
        add      a3, a2, a3
        l32i     a3, a3, 0
        slli     a2, a2, 2            /* multiply by 4 */
        add      a8, a0, a2
        addi     a8, a8, 16
        jx       a3
        /* wb = ipwb; a8 = ar_save_area_ptr + (ipwb+1)*16 */
restore1_28: /* ipwb = 0 */
        l32i     a4, a8, 0
        l32i     a5, a8, 4
        l32i     a6, a8, 8
        l32i     a7, a8, 12
        addi     a12, a8, 16
        rotw     1
restore1_24: /* ipwb = 1 */
        l32i     a4, a8, 0
        l32i     a5, a8, 4
        l32i     a6, a8, 8

```



```

restore2_16: /* ipwb = 4 */
    addi    a8,a4,16
    l32i    a4,a8,0
    l32i    a5,a8,4
    l32i    a6,a8,8
    l32i    a7,a8,12
    rotw    1
restore2_12: /* ipwb = 3 */
    addi    a8,a4,16
    l32i    a4,a8,0
    l32i    a5,a8,4
    l32i    a6,a8,8
    l32i    a7,a8,12
    rotw    1
restore2_8: /* ipwb = 2 */
    addi    a8,a4,16
    l32i    a4,a8,0
    l32i    a5,a8,4
    l32i    a6,a8,8
    l32i    a7,a8,12
    rotw    1
restore2_4: /* ipwb = 1 */
    addi    a8,a4,16
    l32i    a4,a8,0
    l32i    a5,a8,4
    l32i    a6,a8,8
    l32i    a7,a8,12
    rotw    1
restore2_0:
    l32i    a5,a4,20
    l32i    a6,a4,24
    l32i    a7,a4,28
    l32i    a4,a4,16
    rotw    1

    /* set wstart to what the user had */
    wsr     a0,EXCSAVE_0
    l32r    a0, sr_save_area_ptr

    l32i    a0, a0, (WINDOWSTART*4)
    wsr     a0,WINDOWSTART
    rsr     a0, WINDOWBASE
    wsr     a0, WINDOWBASE /* no-op but to avoid iss problem */

    l32r    a0, ar_save_ptr
    s32i    a1, a0, 0 /* save a1, we don't need loc anymore*/

    /* restore ICOUNT & ICOUNTLVL */
    l32r    a0, sr_save_area_ptr
    movi.n  a1, 0
    wsr     a1, ICOUNTLEVEL // first lower icountlevel to 0
    isync
    l32i    a1, a0, (ICOUNT*4)
    wsr     a1, ICOUNT // now write icount.
    isync
    l32i    a1, a0, (ICOUNTLEVEL*4)
    wsr     a1, ICOUNTLEVEL // finally set icountlvl
    isync

    /* Enable WOE */
    //l32r    a1, .LEWOE
    //rsr     a0, PS
    //or      a0, a0, a1
    //wsr     a0, PS
    //rsync
    //l32r    a0, save_area_ptr

    // Put ar_save_area_ptr back into a0 so
    // that we can restore a1
    l32r    a0, ar_save_ptr
    l32i    a1, a0,0

```

```

        rsr    a0,EXCSAVE_0

        rfi    0

        .align 4
_flush_i_cache:
        entry  sp, 48
        dhwb  a2, 0 /* force it out of the data cache (if present) */
        /* Use ihi for a little more efficiency */
        ihi   a2, 0 /* invalidate in i-cache (if present) */
        isync          /* just for safety sake */
        retw.n

```

// Functions to help us out when running inside simulator.

```

        .align 4
_xmon_out:
        entry  sp,16
        mov.n  a3,a2 // pass the 2nd arg as the first arg.
        movi.n a2,-2 // sys_xmon_out
        or     a4,a4,a4 // force window overflow before simcall
        simcall
        retw.n

```

```

        .align 4
_xmon_in:
        entry  sp,16
        movi.n a2,-3
        or     a4,a4,a4 // force window overflow before simcall
        simcall
        retw.n

```

```

        .align 4
_xmon_flush:
        entry  sp,16
        movi.n a2,-4
        or     a4,a4,a4 // force window overflow before simcall
        simcall
        retw.n
        .align 4

```

```

_xmon_init:
        entry  sp,16
        movi.n a2,-7
        or     a4,a4,a4 // force window overflow before simcall
        simcall
        retw.n
        .align 4

```

```

.global _xmon_crash
_xmon_crash:
        entry  sp, 16
        .byte  0,0,0,0,0,0
        retw.n

```

```

# unsigned _xmon_get_cpenable()
#

```

```

        .global _xmon_get_cpenable
        .align 4
_xmon_get_cpenable:
        entry  sp, 16
#ifdef CPENABLE_OFFSET
        rsr    a2, CPENABLE
#endif
        retw

```



```

# void _xmon_set_cpenable(unsigned value)
#
# a2 -- holds the value to set cpenable to
#

        .global _xmon_set_cpenable
        .align 4
_xmon_set_cpenable:
        entry    sp, 16
#ifdef CPENABLE_OFFSET
        wsr      a2, CPENABLE
        rsync
#endif
        retw

# void _xmon_set_user_register(unsigned user_register, unsigned value, unsigned *execute_here)
#
# a2 -- user_register
# a3 -- value
# a4 -- pointer to memory to execute from
#

        .align 4

.wur0_instruction:
        .word    0x00f30000
.wur0_insn_ptr:
        .word    .wur0_instruction

.wur0_placeholder_ptr:
        .word    .wur0_instruction_placeholder

        .align 4
.global _xmon_set_user_register
_xmon_set_user_register:
        entry    sp, 48

# a6 -- temporary for moving memory
# a5 -- pointer to wur0_placeholder
# a4 -- points to the RAM location we will
#       execute from, move the base instruction
#       (including the retw) to that point.

        l32r     a5, .wur0_placeholder_ptr
        l32i     a6, a5, 0
        s32i     a6, a4, 0
        l32i     a6, a5, 4
        s32i     a6, a4, 4

# a5 -- available again, now used to load the
#       base wur instruction which we will now
#       modify for the correct ar and user register
#       number
# a6 -- holds the modified instruction

        l32r     a5, .wur0_insn_ptr
        l32i     a6, a5, 0

# a2 -- holds the user register we are going to write
# a4 -- holds the location in memory that we are going
#       to execute from
# a6 -- holds the instruction we are going to execute

```

```

        slli    a2, a2, 8
        or      a6, a6, a2

# a2 -- Can be used as a temporary now, to
#       OR in the 3, which is the register that
#       holds the value we are going to write to
#       the user register

        movi    a2, 3
        slli    a2, a2, 4
        or      a6, a6, a2

# a2 -- Temporary for merging instructions
# a4 -- pointer to the location we are going to execute
#       from
# a5 -- Holds the value we load from our execution point
# a6 -- The instruction that we are going to execute

        l32i    a5, a4, 0

# Need to merge our 24-bit instruction with 8 bits
# from our execute point
# Want to use the lower 24 bits from a6,
# and the upper 8-bits from a5

        movi    a2, 0xff
        slli    a2, a2, 24
        and     a5, a5, a2
        or      a6, a6, a5
        s32i    a6, a4, 0

# Flush the cache
        mov     a10, a4
        call8   _flush_i_cache

        jx      a4

# Want the upper 24-bits from a6, and the
# lower 8-bits from a4

.align 4
.wur0_instruction_placeholder:
        or      a0, a0, a0
        retw

#
#
# Data for _xmon_get_user_register
#
#

        .align 4
.rur0_insn:
        .word   0x00e30000

.rur0_insn_ptr:
        .word   .rur0_insn

.rur_placeholder_ptr:
        .word   .rur_instruction_placeholder

# unsigned int _xmon_get_user_register(unsigned user_register, unsigned *execute_here)
#

```

```

#      a2(input)  -- user_register
#      a2(output) -- contains the value
#      a3(input)  -- address for executing instructions

        .align 4
.global _xmon_get_user_register
_xmon_get_user_register:
    entry  sp, 48

# a5 -- temporary for moving memory
# a4 -- Points to our rur instruction including ret
#      that we are going to copy to the execution point
# a3 -- Points to the execution point

        l32r    a4, .rur_placeholder_ptr
        l32i    a5, a4, 0
        s32i    a5, a3, 0
        l32i    a5, a4, 4
        s32i    a5, a3, 4

# a4 -- Temp that Points to the rur0 instruction
# a6 -- will hold the rur instruction throughout

        l32r    a4, .rur0_insn_ptr
        l32i    a6, a4, 0

# Shift the user register number to the correct
# offset and OR it into our instruction
# a2 -- Holds the user register being read
# a6 -- instruction being massgaed

        slli    a2, a2, 4
        or      a6, a6, a2

# Now need to set the r-field of the instruction
# to be 2, which is the return value of this function
# a5 -- Temp that holds the constant being ord in
# a6 -- The instruction being massaged

        movi    a5, 2
        slli    a5, a5, 12
        or      a6, a6, a5

# Now load in the word from where we are going to execute
# the rur, merge our rur instruction, and store that word
# back to memory.
# a2 -- Temp for masking
# a3 -- Points to the correct memory location
# a5 -- Holds the WORD we are manipulating

        l32i    a5, a3, 0

# In Little Endian we save the MSB and put our
# instruction in the lower 3 bytes

        movi    a2, 0xff
        slli    a2, a2, 24
        and     a5, a5, a2
        or      a6, a6, a5
        s32i    a6, a3, 0

# Clear the cache line
# a3 -- Points to the location being cleared

        mov     a6, a3
        call4   _flush_i_cache

        movi    a2, 0
        jx      a3

```

```

# A place holder that will be dynamically replaced with
# the correct rur instruction

        .align 4
.rur_instruction_placeholder:
        or      a0, a0, a0
        retw.n

        .global g_dummy_entry_instruction
        .global g_dummy_retw_instruction
        .global g_dummy_entry_ptr
        .global g_dummy_retw_ptr

        .align 4
g_dummy_entry_instruction:
        entry   sp, 16

        .align 4
g_dummy_retw_instruction:
        retw.n

        .align 4
g_dummy_entry_ptr:
        .word   g_dummy_entry_instruction

        .align 4
g_dummy_retw_ptr:
        .word   g_dummy_retw_instruction

# void _xmon_execute_here(unsigned a4_value, void *execute_here);
#
# a2 -- value to be stuffed into a4
# a3 -- execute the instructions at this address
#

        .global _xmon_execute_here
        .align 4
_xmon_execute_here:
        entry   sp, 16

# a8 will be the a4 value after the call4 to the address
        mov     a8, a2
        callx4  a3
        retw

```

II. Xtensa-mon.c

```

/*****
*****
*
*   The following gdb commands are supported:
*
* command          function                      Return value
*
* g                return the value of the CPU registers  hex data or ENN
* G                set the value of the CPU registers      OK or ENN
*
* mAA..AA,LLLL    Read LLLL bytes at address AA..AA       hex data or ENN
* MAA..AA,LLLL:    Write LLLL bytes at address AA..AA      OK or ENN

```

```

*
*   c           Resume at current address           SNN   ( signal NN)
*   CAA..AA     Continue at address AA..AA          SNN
*
*   s           Step one instruction               SNN
*   sAA..AA     Step one instruction from AA..AA    SNN
*
*   k           kill
*
*   ?           What was the last signal ?          SNN   (signal NN)
*
*   bBB..BB     Set baud rate to BB..BB            OK or BNN, then sets
*                                           baud rate
*
* All commands and responses are sent with a packet which includes a
* checksum. A packet consists of
*
* $<packet info>#<checksum>.
*
* where
* <packet info> :: <characters representing the command or response>
* <checksum>    :: < two hex digits computed as modulo 256 sum of <packetinfo>>
*
* When a packet is received, it is first acknowledged with either '+' or '-'.
* '+' indicates a successful transfer. '-' indicates a failed transfer.
*
* Example:
*
* Host:           Reply:
* $m0,10#2a       +$00010203040506070809101112131415#42
*
* *****/

```

```

#include <stdio.h>
#include <signal.h>
#include <machine/specreg.h>
#include <machine/xtl000.h>
#include "DebugExceptionVectorHandler.h"
#include "uart.h"
#include "xtensa-libdb.h"

#define WS_MASK (~(0)<<(NUM_AREGS/4))
#ifdef IS_LITTLE_ENDIAN

#define IS_BREAKN(p) ((p)[0]==0x2d && ((p)[1]&0xf0)==0xf0)
#define IS_BREAK(p) ((p)[2]==0x00 && ((p)[0]&0x0f)==0x00 && ((p)[1]&0xf0)==0x40)
#define BREAKNO(p) (IS_BREAKN(p) ? ((p)[1]&0x0f) : -1 )
#define BREAK_S(p) ((p)[1]&0x0f)
#define BREAK_T(p) (((p)[0]&0xf0)>>4)
#else

#define IS_BREAKN(p) ((p)[0]==0xd2 && ((p)[1]&0x0f)==0x0f)
#define IS_BREAK(p) ((p)[2]==0x00 && ((p)[0]&0xf0)==0x00 && ((p)[1]&0xf0)==0x04)
#define BREAKNO(p) (IS_BREAKN(p) ? (((p)[1]&0xf0)>>4) : -1 )
#define BREAK_S(p) (((p)[1]&0xf0)>>4)
#define BREAK_T(p) ((p)[0]&0xf0)

#endif

#define SR_REG(n) (_sr_registers[ (n) ])

/* Macros to extract fields of PS */
#define GET_PSINTLVL(ps) ((ps)&0xf)
#define GET_PSUSRMODE(ps) (((ps)>>5)&0x1)
#define GET_PSOWB(ps) (((ps)>>8)&0xf)
#define GET_PSCALLINC(ps) (((ps)>>16)&0x3)
#define GET_PSWOE(ps) (((ps)>>18)&0x1)

/* Imported functions */
extern void _flush_i_cache( char *);
extern int _xmon_out(char c);

```

[illegible]

```

/* Parameters:
   We'll need to create a configuration process that generates
   many of these defines.
*/
/*****
/* BUFMAX defines the maximum number of characters in inbound/outbound buffers*/
/* at least NUMREGBYTES*2 are needed for register packets */
#define BUFMAX 2048

#define AR_SAVE_SIZE (4*NUM_AREGS)
#define SR_SAVE_SIZE (4*256)

#define PC EPC_0
#define DBG_EPS EPS_0

long _ar_registers[AR_SAVE_SIZE/(sizeof (long))];
long _sr_registers[SR_SAVE_SIZE/(sizeof (long))];

/* !0 means we are running on the board, defined by linker */
extern void *IN_SIMULATOR;
int _in_simulator = (int)&IN_SIMULATOR;
int _initialized = 0; /* !0 means we've been initialized */

static const char hexchars[]="0123456789abcdef";

/* string functions */
int _strlen( char *cp )
{
    int i;
    if( cp == 0 )
        return 0;
    i = 0;
    while ( *cp )
    {
        i++;
        cp++;
    }
    return i;
}

char *_strcpy( char *d, char *s )
{
    char *cp = d;
    if( d && s )
    {
        while( *s )
        {
            *cp = *s;
            cp++;
            s++;
        }
        *cp = 0;
    }
    return d;
}

void _memset(unsigned char *ptr, unsigned char value, int num)
{
    while (num > 0)
    {
        *ptr = value;
        ++ptr;
        --num;
    }
}

```

0060013E:100300

```

/* Log errors for transmission */
#define LOG_SIZE 100
static char *error_log_end;
static char error_log[LOG_SIZE];

static void mon_error_clear()
{
    error_log_end = error_log;
    error_log_end[0] = 0;
}

static void mon_error( char *msg )
{
    if( error_log_end == 0 )
        error_log_end = error_log;
    while( *msg )
    {
        if( error_log_end < &error_log[LOG_SIZE-1] )
            *error_log_end++ = *msg++;
        else
            break;
    }
    error_log_end[0] = 0;
}

/* Convert ch from a hex digit to an int */
static int
hex(ch)
    unsigned char ch;
{
    if (ch >= 'a' && ch <= 'f')
        return ch-'a'+10;
    if (ch >= '0' && ch <= '9')
        return ch-'0';
    if (ch >= 'A' && ch <= 'F')
        return ch-'A'+10;
    return -1;
}

/* scan for the sequence $<data>#<checksum> */
static void
getpacket(buffer)
    char *buffer;
{
    unsigned char checksum;
    unsigned char xmitcsum;
    int i;
    int count;
    unsigned char ch;

    do
    {
        /* wait around for the start character, ignore all other characters */
        while ((ch = (getDebugChar() & 0x7f)) != '$') ;

        checksum = 0;
        xmitcsum = -1;

        count = 0;

        /* now, read until a # or end of buffer is found */
        while (count < BUFMAX)
        {
            ch = getDebugChar() & 0x7f;
            if (ch == '#')

```



```

        break;
        checksum = checksum + ch;
        buffer[count] = ch;
        count = count + 1;
    }

    if (count >= BUFMAX)
        continue;

    buffer[count] = 0;

    if (ch == '#')
    {
        xmitcsum = hex(getDebugChar() & 0x7f) << 4;
        xmitcsum |= hex(getDebugChar() & 0x7f);
    }

    /* Humans shouldn't have to figure out checksums to type to it. */
    putDebugChar('+');
    return;

#endif

    if (checksum != xmitcsum)
        putDebugChar('-'); /* failed checksum */
    else
    {
        putDebugChar('+'); /* successful transfer */
        /* if a sequence char is present, reply the sequence ID */
        if (buffer[2] == ':')
        {
            putDebugChar(buffer[0]);
            putDebugChar(buffer[1]);
            /* remove sequence chars from buffer */
            count = _strlen(buffer);
            for (i=3; i <= count; i++)
                buffer[i-3] = buffer[i];
        }
        flushDebug();
    }
}

while (checksum != xmitcsum);
}

/* send the packet in buffer. */

/* Convert the memory pointed to by mem into hex, placing result in buf.
 * Return a pointer to the last char put in buf (null), in case of mem fault,
 * return 0.
 * If MAY_FAULT is non-zero, then we will handle memory faults by returning
 * a 0, else treat a fault like any other fault in the stub.
 */

static unsigned char *
mem2hex(mem, buf, count, may_fault)
    unsigned char *mem;
    unsigned char *buf;
    int count;
    int may_fault;
{
    unsigned char ch;
    while (count-- > 0)
    {
        ch = *mem++;
        *buf++ = hexchars[ch >> 4];
        *buf++ = hexchars[ch & 0xf];
    }

    *buf = 0;

    return buf;
}

```

```

static void
putpacket(buffer)
    unsigned char *buffer;
{
    unsigned char checksum;
    unsigned char ack;
    int count;
    unsigned char ch;

    /* $<packet info>#<checksum>. */
    do
    {
        putDebugChar('$');
        checksum = 0;
        count = 0;

        while (ch = buffer[count])
        {
            putDebugChar(ch);
            checksum += ch;
            count += 1;
        }

        putDebugChar('#');
        putDebugChar(hexchars[checksum >> 4]);
        putDebugChar(hexchars[checksum & 0xf]);

        flushDebug();
        ack = getDebugChar();
        ack = ack & 0x7f;

//        led_display_ok();
        /*
        if (ack != '+')
        {
            //        char buf[8];
            //        _memset(buf, 0, sizeof(buf));
            //        putDebugString("---");
            //        mem2hex(&ack, buf, 1, 0);
            putDebugString(buf);
            putDebugString("---");
        }
        else
            putDebugChar('Y');
        */
    }
    while (ack != '+');
}

static char remcomInBuffer[BUFMAX];
static char remcomOutBuffer[BUFMAX];

static unsigned char g_execute_here[1024];

static void bad_protocol()
{
    _strcpy( remcomOutBuffer, "Error: garbled command" );
}

static void aok()
{
    _strcpy( remcomOutBuffer, "OK" );
}

/* Decode a hex string and write it into memory. */
static char *
write_mem(buf, mem, count, flush, verify)
    register unsigned char *buf;
    register unsigned char *mem;

```

```

    int count;
    int flush;
    int verify;
{
    int i;
    unsigned char ch;
    unsigned char *start = mem;

    for (i=0; i<count; i++)
    {
        ch = hex(*buf++) << 4;
        ch |= hex(*buf++);
        *mem = ch;
        if( verify && *mem != ch )
            return 0;
        mem += 1;
    }

    if( flush ) {
        while( count >= 0 ) {
            _flush_i_cache( start );
            count -= 4;
            start += 4;
        }
        /* we do one more flush just in case the last
           instruction straddled two cached line */
        _flush_i_cache( start );
    }

    return mem;
}

/*
 * While we find nice hex chars, build an int.
 * Return number of chars processed.
 */

static int
hexToInt(char **ptr, int *intValue)
{
    int numChars = 0;
    int hexValue;

    *intValue = 0;

    while (**ptr)
    {
        hexValue = hex(**ptr);
        if (hexValue < 0)
            break;

        *intValue = (*intValue << 4) | hexValue;
        numChars ++;

        (*ptr)++;
    }

    return (numChars);
}

static void set_icount_for_single_step(int intlevel)
{
    /* set the icount level to one more than the interrupt level,
       This will allow single-stepping through handlers */
    SR_REG(ICOUNT) = -2;
    SR_REG(ICOUNTLEVEL) = intlevel < DEBUG_INTERRUPT_LEVEL ? intlevel+1 :
        DEBUG_INTERRUPT_LEVEL;
}

```

```

/* The _handle_exception function is best modeled as a state machine */
static int state; /* zero'ed by during bss initialization */
/* Thus XMON_INITIAL must be zero. */
#define XMON_INITIAL 0 /* start up xmon */
#define XMON_CONTROL 1 /* xmon is stopped, polling
                        commands from host */
#define XMON_RUNNING 2 /* xmon is running, waiting for
                        an external event */
#define XMON_RESUMING 3 /* an interrupt, not the serial
                        interrupt, has occurred. */

/* Data structe to keep track of hw breakpoints */
struct hw_break_info {
    int free;
    char *addr;
    int reg_number;
} hw_break[NIBREAK] = HW_BREAK_INIT;

/* special breaks are how we detect SIGINT and SIGILL */
typedef void (*special_breakpoint_handler)(int *,int *);
static void sigint_handler(int *,int*);
static void _tell_gdb(int);

struct special_breakpoint {
    char *address;
    special_breakpoint_handler f;
    char saved_inst[3];
} special_break[] = {
    { (char *)UART_VECTOR, sigint_handler },
#ifdef UART_SECOND_VECTOR
    { (char *)UART_VECTOR_2, sigint_handler },
#endif
    { (char *)0, (special_breakpoint_handler)0 }
};

static void init_special_breaks()
{
    /* nothing to do right now */
}

static void set_special_breaks()
{
    struct special_breakpoint *b;
    b = special_break;
    while( b->address )
    {
#ifdef ISAUSEDENSITYINSTRUCTION
        b->saved_inst[0] = b->address[0];
        b->saved_inst[1] = b->address[1];
#ifdef IS_LITTLE_ENDIAN
        b->address[0] = 0x2d;
        b->address[1] = 0xf1;
#else
        b->address[0] = 0xd2;
        b->address[1] = 0x1f;
#endif
#endif
        _flush_i_cache(b->address);
        _flush_i_cache(b->address+1);
#ifdef IS_LITTLE_ENDIAN
        b->address[0] = 0x10;
        b->address[1] = 0x40;
        b->address[2] = 0x00;
#else
        b->address[0] = 0x01;
        b->address[1] = 0x04;
        b->address[2] = 0x00;
#endif
    }
}

```

```
#endif
    _flush_i_cache(b->address);
    _flush_i_cache(b->address+1);
    _flush_i_cache(b->address+2);
#endif
    b++;
}
}

static void clear_special_breaks()
{
    struct special_breakpoint *b;
    b = special_break;
    while( b->address )
    {
#ifdef ISAUSEDENSITYINSTRUCTION
        b->address[0] = b->saved_inst[0];
        b->address[1] = b->saved_inst[1];
        _flush_i_cache(b->address);
        _flush_i_cache(b->address+1);
#else
        b->address[0] = b->saved_inst[0];
        b->address[1] = b->saved_inst[1];
        b->address[2] = b->saved_inst[2];
        _flush_i_cache(b->address);
        _flush_i_cache(b->address+1);
        _flush_i_cache(b->address+2);
#endif
        b++;
    }
}

// !!@

static void do_special_breaks(int *state, int *sigval)
{
    char *pc;
    struct special_breakpoint *b;

    b = special_break;
    pc = (char *)SR_REG(PC);

    while( b->address )
    {
        if( b->address == pc && b->f )
        {
            b->f(state, sigval);
            return;
        }
        b++;
    }

    *state = XMON_CONTROL;
    *sigval = SIGTRAP;
}

#if 0
void set_ps( int eps )
{
    REG(PSINTLVL)   = GET_PSINTLVL(eps);
    REG(PSUSRMODE)  = GET_PSUSRMODE(eps);
    REG(PSOWB)      = GET_PSOWB(eps);
    REG(PSCALLINC)  = GET_PSCALLINC(eps);
    REG(PSWOE)      = GET_PSWOE(eps);
}
#endif
```

```

#if 0
void flash_value(unsigned int value)
{
    int i = 0;

    for(i = 28; i >= 0; i=i-4)
    {
        int number = (value >> i) & 0xf;

        led_blank();
        led_pause(100000);
        led_display_digit(number);
        led_pause(100000);
    }
}
#endif

void setup_ps()
{
    /* the code to set up the PS works quite differently depending
       on whether or not the uart is on interrupt level one. */
    #if UART_INTERRUPT_LEVEL == 1
    {
        unsigned int real_ps = 0;
        real_ps = SR_REG(EPS_0);

        // We can figure out if PS.UM was set by looking
        // at which vector we came from
        // UART_VECTOR is actually the UserExceptionVector

        if ( SR_REG(UART_EPC) == UART_VECTOR )
        {
            // Since we are coming from UserExceptionVector
            // turn on the PS.UM mode bit.
            real_ps = real_ps | (1 << 4);

            // Assume that WOE is always enabled for user code.
            real_ps = real_ps | (1 << 17);
        }
        else
        {
            // We are coming from the KernelExceptionVector
            // So we leave PS.UM disabled, and we take
            // WOE from the current PS.
            real_ps = real_ps | ( SR_REG(EPS_0) & (1 << 17) );
        }

        // Set interrupt level to 0
        real_ps = real_ps & ~0xf;
        SR_REG(EPS_0) = real_ps;
    }
    #elif UART_INTERRUPT_LEVEL != -1
        SR_REG(EPS_0) = SR_REG(UART_EPS);
    #endif
}

/* If we see a break on the UART then we simulate a SIGINT */
static void sigint_handler( int *state, int *sigval )
{
    int interrupts = SR_REG(INTERRUPT);
    int c;

    // led_display_digit(7);
    // led_pause(100000);

    // flash_value(interrupts);

```

```

    if( interrupts & UART_INTERRUPT )
    {
        //      led_display_digit(8);
        //      led_pause(100000);

        /* the received interrupt was the serial interrupt for the
           port that is being used for xmon control. GDB has
           requested that xmon break. */
        *state = XMON_CONTROL;
        c = getDebugChar();      /* read the break char, to avoid
                                   gdb protocol sync problems */
        *sigval = SIGINT;        /* indicate SIGINT */

        /* unwind the interrupt: set up registers so that it appears
           we returned from interrupt handler.
        */

        // !!@
        //      led_display_digit(9);
        //      led_pause(100000);

#ifdef UART_INTERRUPT_LEVEL != -1
        SR_REG(PC) = SR_REG(UART_EPC);
#endif
        setup_ps();
    }
    else
    {
        /* this is some other level2 interrupt. Special breaks were
           already cleared so set the state into resume mode and
           set the icount up for a single instruction. This will
           allow us to step over the instruction and restore the
           break.*/
        //      led_display_digit(10);
        //      led_pause(100000);

        *state = XMON_RESUMING;
        set_icount_for_single_step( SR_REG(DBG_EPS) & 0x0f );
    }
}

void
_handle_exception()
{
    int n;
    int sigval = SIGTRAP;
    unsigned char *pc;
    pc = (unsigned char *)SR_REG(PC);

    /* reset icount, so that we can continue, in case
       we came here because of an icount interrupt */
    SR_REG(ICOUNT) = 0;
    SR_REG(ICOUNTLEVEL) = 0;

    /* when we return enable all interrupts except timers */
#ifdef TIMER_INTERRUPT_MASK
    //      SR_REG(INTENABLE) = ALL_INTERRUPT_MASK & (~TIMER_INTERRUPT_MASK);
    SR_REG(INTENABLE) = ALL_INTERRUPT_MASK;
#else
    SR_REG(INTENABLE) = ALL_INTERRUPT_MASK;
#endif

    for(;;)
    {
        switch( state )
        {
            case XMON_INITIAL:
                if( !_in_simulator )
                {

```

```

        /* We're running on the board, so flash the status
        LEDs a few times */
        _uart_init((uart_dev_t *)XT1000_DUART_1_ADDR, B38400);
        _uart_enable_rcvr_int((uart_dev_t *)XT1000_DUART_1_ADDR);
        led_display_ok();
        led_display_digit(2);
//    }
    else
    {
        _xmon_init();
    }
    putDebugString("XMON R2.5 ");
    putDebugString(_in_simulator ? " running on iss\n"
        : " running on eval board\r\n");

    _initialized = 1;

    init_special_breaks();
    state = XMON_CONTROL;
    continue;

case XMON_CONTROL:    /* let host control us */
    _tell_gdb( sigval );
    set_special_breaks();
    state = XMON_RUNNING;
    return;           /* return to user program */

case XMON_RUNNING:
    if(IS_BREAK(pc)) {
        unsigned s = BREAK_S(pc);
        unsigned t = BREAK_T(pc);
        if( s==1 && (t <= 1)) {
            switch(SR_REG(EXCCAUSE)) {
                case EXCCAUSE_ILLEGAL:
                    sigval = SIGILL;
                    break;
                case EXCCAUSE_SYSCALL:
                    sigval = SIGTRAP;
                    break;
                case EXCCAUSE_IFETCHERROR:
                case EXCCAUSE_LOADSTOREERROR:
                    sigval = SIGSEGV;
                    break;
                case EXCCAUSE_LEVEL1INTERRUPT:
                    sigval = SIGINT;
                    break;
            }
            clear_special_breaks();
            state = XMON_CONTROL;

            /* pretend as though we caught it
            at the point of occurrence */
            // !!@
            SR_REG(PC) = SR_REG(EPC_1);
            setup_ps();
            continue;
        }
    }
    n = BREAKNO(pc);
    switch(n)
    {
    default:
        clear_special_breaks();
        state = XMON_CONTROL;
        break;

    case 1:           /* special breakpoint */
        clear_special_breaks();

        /* keep in mind that the state can be changed by the
        do_special_breaks handler. */

```

00000000 " 9280486 1003000


```

        do_special_breaks(&state, &sigval);
        break;
    }
    continue;

case XMON_RESUMING:
    /* we have taken an interrupt that was not the serial
       interrupt and have now executed the original instruction
       at the interrupt vector. Now restore the break instruction
       so that interrupts continue to work and resume. */

    set_special_breaks();
    state = XMON_RUNNING;
    return;
}
}

/*
 * This function does all command processing for interfacing to gdb.
 */
unsigned char dummy[4];
unsigned int  user_register_value;

unsigned int execution_space[2];

static unsigned char *
get_reg_ptr(const unsigned int libdb_target_number)
{
    unsigned char *reg_ptr = NULL;
    unsigned      old_cpe = 0;
    int           offset = 0;

    switch( GET_TARGET_REG_TYPE(libdb_target_number) )
    {
        case REGTYPE_AR:
            offset = GET_TARGET_REG_INDEX(libdb_target_number);
            reg_ptr = (unsigned char *)&_ar_registers[offset];
            break;

        case REGTYPE_SPECIAL_REG:
            offset = GET_TARGET_REG_INDEX(libdb_target_number);
            reg_ptr = (unsigned char *)&_sr_registers[offset];
            break;

        case REGTYPE_USER_REG:
            old_cpe = _xmon_get_cpenable();
            _xmon_set_cpenable((unsigned)-1);
            user_register_value = _xmon_get_user_register( GET_TARGET_REG_INDEX(libdb_target_number),
                                                         &execution_space[0] );
            _xmon_set_cpenable(old_cpe);
            reg_ptr = (unsigned char *)&user_register_value;
            break;

        default:
            reg_ptr = NULL;
            break;
    }

    return reg_ptr;
}

static unsigned int
set_reg_value(const unsigned int libdb_target_number, const unsigned int value)
{
    int success = 0;

    if ( GET_TARGET_REG_TYPE(libdb_target_number) == REGTYPE_USER_REG )
    {

```



```

g_execute_here[1] = g_dummy_entry_ptr[1];
g_execute_here[2] = g_dummy_entry_ptr[2];

index = 3;

_flush_i_cache( (char *)&g_execute_here[0] );

while (pInstruction && *pInstruction)
{
    // Skip the ':' characters

    ++pInstruction;
    converted = hexToInt(&pInstruction, &dummy);
    if (converted != 2)
        goto exit_gracefully;

    g_execute_here[index] = (unsigned char)dummy;
    _flush_i_cache( (char *)&g_execute_here[index] );
    ++index;
}

g_execute_here[index++] = g_dummy_retw_ptr[0];
_flush_i_cache( (char *)&g_execute_here[index] );

g_execute_here[index++] = g_dummy_retw_ptr[1];
_flush_i_cache( (char *)&g_execute_here[index] );

g_execute_here[index++] = g_dummy_retw_ptr[2];
_flush_i_cache( (char *)&g_execute_here[index] );

_xmon_execute_here(a4_value, &g_execute_here[0]);

success = 1;

exit_gracefully:

_xmon_set_cpenable( old_cpe );

return success;
}

static void
_tell_gdb ( int sigval )
{
    int tt;
    int addr;
    int length;
    int value;
    int intlevel;
    int woe;
    char *ptr;
    unsigned long *sp;
    unsigned int wb, pc;
    int i;
    struct hw_break_info *bp;

    wb = SR_REG(WINDOWBASE);
    sp = (unsigned long *) reg_at_wb( SP_REGNUM, wb );
    pc = SR_REG(PC);
    intlevel = SR_REG(DBG_EPS) & 0x0f;
    woe = SR_REG(DBG_EPS) & 0x040000;

    /* If we find that window overflow/underflow enabled and
       the interrupt level zero then we can safely save all
       the registers to the stack.
    */
}

```

```

if( woe != 0 && intlevel==0 ) {
    if(save_to_stack()!=0) {
        ptr = remcomOutBuffer;
        *ptr++ = 'E';
        if( error_log_end != 0 && error_log_end < &error_log[LOG_SIZE-1]) {
            error_log_end[0] = 0;
            _strcpy( ptr, error_log );
            mon_error_clear();
        } else
            _strcpy( ptr, "Error in save_to_stack\n" );
        putpacket(remcomOutBuffer);
    }
}

ptr = remcomOutBuffer;

/* Tell gdb that we have stopped */
*ptr++ = 'S';
*ptr++ = hexchars[signal >> 4];
*ptr++ = hexchars[signal & 0xf];
*ptr++ = 0;
putpacket(remcomOutBuffer);

while (1)
{
    remcomOutBuffer[0] = 0;

    getpacket(remcomInBuffer);
    switch (remcomInBuffer[0])
    {
        case '?':
            remcomOutBuffer[0] = 'S';
            remcomOutBuffer[1] = hexchars[signal >> 4];
            remcomOutBuffer[2] = hexchars[signal & 0xf];
            remcomOutBuffer[3] = 0;
            break;

        case 'E': /* xtensa specific */
            /* send the error message log back */
            /* mem2hex( error_log, remcomOutBuffer, error_log_end-error_log, 0 );*/
            break;

        case 'd':
            /* toggle debug flag */
            break;

        #if 0
        case 'g':
            /* return the value of the CPU registers */
            {
                ptr = mem2hex( (char *)_registers, remcomOutBuffer,
                    SAVE_AREA_SIZE, 0 );
            }
            break;

        case 'G':
            /* set the value of the CPU registers - return OK */
            {
                /* We allow the user to set any registers without checking.
                 Users can wedge the board if they load inconsistent values
                 into the registers */
                write_mem( &remcomInBuffer[1], (char *)_registers,
                    SAVE_AREA_SIZE, 0, 0 );
                ack();
            }
            break;

        #endif

        case 'p':
            ptr = &remcomInBuffer[1];
            if( hexToInt(&ptr, &addr) )
            {
                unsigned char *reg_ptr = NULL;

```

```

        // !!@
        //if(mem2hex( (char *)&_registers[addr], remcomOutBuffer, 4, 0 ))
        //    break;
        reg_ptr = get_reg_ptr( addr);
        if (reg_ptr != NULL)
        {
            if(mem2hex( reg_ptr, remcomOutBuffer, 4, 0 ))
                break;
        }
        _strcpy( remcomOutBuffer, "Error: read failure" );
    }
    else
        bad_protocol();
    break;

case 'P':
    ptr = &remcomInBuffer[1];
    if( hexToInt(&ptr, &addr) && *ptr++ == '='
        && write_mem( ptr, (char *)&value, 4, 0, 0 ))
    {
        unsigned char *reg_ptr = NULL;
        if (set_reg_value(addr, value))
        {
            aok();
        }
        else
        {
            _strcpy( remcomOutBuffer, "Error: write failure" );
        }
        // !!@
        //    _registers[addr] = value;
    }
    else
        bad_protocol();
    break;

case 'M': /* MAA..AA,LLLL Read LLLL bytes at address AA..AA */
    /* Try to read %x,%x. */

    ptr = &remcomInBuffer[1];

    if (hexToInt(&ptr, &addr)
        && *ptr++ == ','
        && hexToInt(&ptr, &length))
    {
        if (mem2hex((char *)addr, remcomOutBuffer, length, 1))
            break;
        _strcpy( remcomOutBuffer, "Error: read failure" );
    }
    else
        bad_protocol();
    break;

case 'M': /* MAA..AA,LLLL: Write LLLL bytes at address AA.AA return OK */
    /* Try to read '%x,%x:'. */

    ptr = &remcomInBuffer[1];

    if (hexToInt(&ptr, &addr)
        && *ptr++ == ','
        && hexToInt(&ptr, &length)
        && *ptr++ == ':')
    {
        if (write_mem(ptr, (char *)addr, length, 1, 1))
            aok();
        else
            _strcpy(remcomOutBuffer, "Error: write failure");
    }
    else
        bad_protocol();

```

```

break;

case 'c':    /* cAA..AA Continue at address AA..AA(optional) */
/* try to read optional parameter, pc unchanged if no parm */

    ptr = &remcomInBuffer[1];
    if (hexToInt(&ptr, &addr))
    {
        SR_REG(PC) = addr;
    }
    // else
    // bad_protocol();
    return;

case 's':
/* use icount mechanism to step a single instruction */
    set_icount_for_single_step(intlevel);
    return;

/* kill the program */
case 'k' :    /* do nothing */
    break;

/* Xtensa specific commands */
case 'X':
    switch( remcomInBuffer[1] )
    {
        case 'q':
            switch (remcomInBuffer[2])
            {
                case 'n':
                    _strcpy( remcomOutBuffer, "XMON2.5" );
                    break;

                case 'p':
                    _strcpy( remcomOutBuffer, "n");
                    break;

                case 'P':
                    _strcpy( remcomOutBuffer, "n");
                    break;

                default:
                    break;
            }
            break;

        case 'e':
            if ( remcomInBuffer[2] == 'x' && remcomInBuffer[3] == 'e' )
            {
                ExecuteSomeInstruction( &remcomInBuffer[4] );
                remcomOutBuffer[0] = '\0';
            }
            break;

        case 'B':
            /* Set a breakpoint using the ibreak registers */
#ifdef NIBREAK==0
            _strcpy( remcomOutBuffer, "Error: configuration has no IBREAK registers");
#else
            ptr = &remcomInBuffer[4];
            if( !hexToInt(&ptr, &addr ) ) {
                bad_protocol();
                break;
            }

            switch( remcomInBuffer[2] ) {
                case 's':    /* set */
                    for( i = 0, bp = hw_break; i < NIBREAK; i++, bp++ )

```

```

        if( bp->free ) {
            bp->free = 0;
            bp->addr = (char *)addr;
            // !!@
            //_registers[bp->reg_number] = addr;
            SR_REG(IBREAKENABLE) |= (1<<i);
            aok();
            break;
        }
        if( i >= NIBREAK )
            _strcpy( remcomOutBuffer, "Error: out of ibreak registers");
        break;
    case 'r':
        /* remove */
        for( i = 0, bp = hw_break; i < NIBREAK; i++, bp++ )
            if( !bp->free && bp->addr == (char *)addr ) {
                bp->free = 1;
                bp->addr = 0;
                // !!@
                //_registers[bp->reg_number] = 0;
                SR_REG(IBREAKENABLE) &= ~(1<<i);
                aok();
                break;
            }
        if( i >= NIBREAK )
            _strcpy( remcomOutBuffer, "Error: breakpoint not found");
        break;
    default:
        break;
    }
#endif
    }
    break;
#endif

/* Test feature */
case 't':
    asm (" std %f30,[%sp]");
    break;
/* Reset */
case 'r':
    asm ("call 0
        nop ");
    break;
#endif

/* Disabled until we can unscrew this properly */
case 'b':
    /* bBB... Set baud rate to BB... */
    {
        int baudrate;
        extern void set_timer_3();

        ptr = &remcomInBuffer[1];
        if (!hexToInt(&ptr, &baudrate))
        {
            _strcpy(remcomOutBuffer,"B01");
            break;
        }

        /* Convert baud rate to uart clock divider */
        switch (baudrate)
        {
            case 38400:
                baudrate = 16;
                break;
            case 19200:
                baudrate = 33;
                break;
            case 9600:
                baudrate = 65;
                break;
            default:
                _strcpy(remcomOutBuffer,"B02");
        }
    }

```



```

/*****
Xtensa hardware-dependent utilities
*****/

/* retrieved register at a particular window base */
static long reg_at_wb( unsigned int reg, unsigned int wb )
{
    unsigned int relocated_reg;
    if( (NUM_AREGS/4) <= wb )
        mon_error( "invalid window base in reg_at_wb\n" );
    if( NUM_VISIBLE_AREGS <= reg )
        mon_error( "invalid register in reg_at_wb\n" );
    relocated_reg = (reg + (wb<<2)) & AREGS_MASK;
    // relocated_reg += ARO_OFFSET/4;

    return _ar_registers[relocated_reg];
}

/* retrieved register in window of interrupt process */
static int reg_at_ipwb( unsigned int reg )
{
    return reg_at_wb( reg, SR_REG(WINDOWBASE) );
}

static int save_to_stack()
{
    int i;
    int ws = SR_REG(WINDOWSTART);
    int wb = SR_REG(WINDOWBASE);
    int callee_win, win;
    long *sp, *caller_sp;

    /* rotate so that first bit of ws corresponds to
       wb+1 */
    ws = (ws >> (wb+1)) | (ws << (NUM_AREGS/4-(wb+1)));
    ws &= WS_MASK;

    /* find first window after ipwb */
    if( ws == 0 )
        mon_error( "window start zero in save_to_stack\n" );
    for( i = 0; (ws&1)==0; i++ )
        ws >>= 1;
    ws >>= 1;
    i++;
    while( ws != 0 )
    {
        win = (wb+i) & WB_MASK;
        if( ws & 1 )
        {
            callee_win = (win+1) & WB_MASK;
            sp = (long *)reg_at_wb( 1, callee_win) - 4;
            sp[0] = reg_at_wb( 0, win );
            sp[1] = reg_at_wb( 1, win );
            sp[2] = reg_at_wb( 2, win );
            sp[3] = reg_at_wb( 3, win );
            i = i+1;
            ws >>= 1;
            continue;
        }
        if( ws & 2 )
        {
            callee_win = (win+2) & WB_MASK;
            sp = (long *)reg_at_wb( 1, callee_win) - 4;
            sp[0] = reg_at_wb( 0, win );
            sp[1] = (long) caller_sp = (long *)reg_at_wb( 1, win );
            sp[2] = reg_at_wb( 2, win );
            sp[3] = reg_at_wb( 3, win );
        }
    }
}

```


Parameter	Unit	Value	Unit	Value
Initial temperature	°C	25	Initial temperature	°C
Final temperature	°C	100	Final temperature	°C
Heating rate	°C/min	10	Heating rate	°C/min
Sample weight	g	0.5	Sample weight	g
Sample size	mm	10 × 10 × 2	Sample size	mm
Sample density	g/cm ³	1.2	Sample density	g/cm ³
Sample purity	%	100	Sample purity	%
Sample origin		Commercial	Sample origin	
Sample grade		High purity	Sample grade	
Sample history		None	Sample history	
Sample storage		Room temperature	Sample storage	
Sample handling		Standard	Sample handling	
Sample analysis		Standard	Sample analysis	
Sample results		Standard	Sample results	
Sample conclusion		Standard	Sample conclusion	

Parameter	Unit	Value	Unit	Value
Initial temperature	°C	25	Initial temperature	°C
Final temperature	°C	100	Final temperature	°C
Heating rate	°C/min	10	Heating rate	°C/min
Sample weight	g	0.5	Sample weight	g
Sample size	mm	10 × 10 × 2	Sample size	mm
Sample density	g/cm ³	1.2	Sample density	g/cm ³
Sample purity	%	100	Sample purity	%
Sample origin		Commercial	Sample origin	
Sample grade		High purity	Sample grade	
Sample history		None	Sample history	
Sample storage		Room temperature	Sample storage	
Sample handling		Standard	Sample handling	
Sample analysis		Standard	Sample analysis	
Sample results		Standard	Sample results	
Sample conclusion		Standard	Sample conclusion	